

Sleep Analytics and Online Selective Anomaly Detection

Tahereh Babaie
School of IT
University of Sydney
Sydney, NSW, Australia
tahereh.babaie@sydney.edu.au

Sanjay Chawla
School of IT
University of Sydney
Sydney, NSW, Australia
sanjay.chawla@sydney.edu.au

Romesh Abeysuriya
School of Physics
University of Sydney
Sydney, NSW, Australia
r.abeyesuriya@physics.usyd.edu.au

ABSTRACT

We introduce a new problem, the Online Selective Anomaly Detection (OSAD), to model a specific scenario emerging from research in sleep science. Scientists have segmented sleep into several stages and stage two is characterized by two patterns (or anomalies) in the EEG time series recorded on sleep subjects. These two patterns are sleep spindle (SS) and K-complex. The OSAD problem was introduced to design a residual system, where all anomalies (known and unknown) are detected but the system only triggers an alarm when non-SS anomalies appear. The solution of the OSAD problem required us to combine techniques from both machine learning and control theory. Experiments on data from real subjects attest to the effectiveness of our approach.

General Terms

Anomaly/novelty detection, Mining rich data types

Keywords

Sleep EEG Anomalies, Dynamic Residue Model

1. INTRODUCTION

Research in human sleep condition has emerged as a rapidly growing area within medicine, biology and physics. A defining aspect of sleep research is the large amount of data that is generated in a typical sleep experiment.

A sleep experiment consists of a human subject, in a state of sleep, whose neural activity is being recorded with Electroencephalography (EEG) [19, 4]. A typical full night EEG time-series, recorded between 4-64 locations on the scalp, at 200 Hz, for eight hours, will generate approximately 300MB of data. A typical clinical study will have between ten and fifty subjects. Surprisingly vast majority of sleep clinics still use a manual process to analyze the recorded EEG time-series. Hence there is considerable interest in automating the analysis of EEG generated from sleep experiments. Scientists have segmented sleep into several stages based on

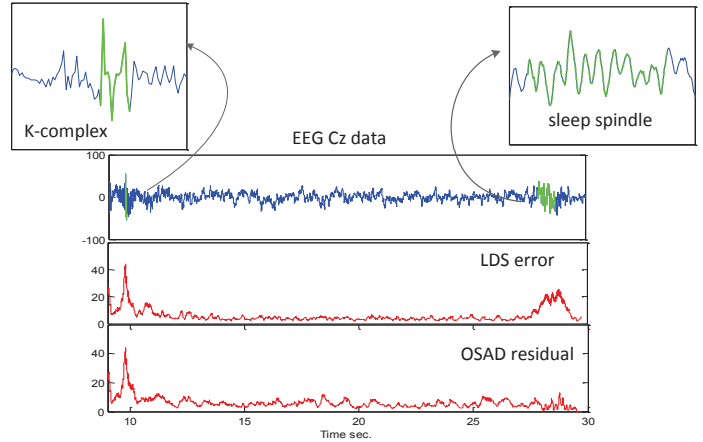


Figure 1: Sleep spindles (SS) along with K-Complexes (KC) are defining characteristics of stage 2 sleep. Both SS and KC will show up as residuals in an LDS system. The OSAD problem will lead to a *new* residual time-series where SS will be automatically suppressed but KC will remain unaffected. Due to relatively high frequency of SS, there are certain situations where sleep scientists only want to be alerted when a non-SS anomaly occurs

the responsiveness of the subject and other physiological features. Of particular important is what is termed as stage 2 (moderately deep sleep). This stage is characterized by two phenomenon that occur in the EEG time series. These are *sleep spindles*, which are transient bursts of neural activity with a characteristic frequency of 12–14 Hz, and *K-Complexes*, which are short, large-amplitude voltage spikes. Both phenomena are implicated in memory consolidation and learning, but the physiology and mechanisms by which they occur are not yet fully understood, see [4, 9, 19, 8].

In order to study these phenomena, they anomalies must be first located and identified in the EEG data. This can be challenging because they occur for an extremely short duration and irregularly. For example, sleep spindles and K-Complexes typically last less than 1 s, and there are only on the order of 100 of these events over the course of an entire night. Identification of these events is further complicated by the presence of artifacts in the data, often caused by movement of the subject, but which can also occur due to electrical noise or loose electrodes connections. These artifacts must be ignored when attempting to identify sleep spindles and K-Complexes. Because the electric fields pro-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WOODSTOCK '97 El Paso, Texas USA

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

duced by the brain are quite weak (the induced electrical potential is on the order of $50\mu\text{V}$), the signals also contain a significant noise component.

In this paper we introduce the Online Selective Anomaly Detection (OSAD) problem which captures a particular scenario in sleep research. As noted above, around 100 sleep spindles will occur during the course of a night. The number of K-Complexes is much fewer. For some experiments scientists are interested in identifying both sleep spindles and K-Complexes but only want to be notified with an alert when a non-spindle anomaly occurs (for example K-Complexes).

The solution of the OSAD problem combines techniques from both data mining and control theory. Data Mining is used to model and infer the normal EEG pattern per subject. Experiments have shown that model parameters do not transfer accurately across to other subjects. In our case we will use a Linear Dynamical System (LDS) to model the EEG time series. Then based on frequency analysis, we infer the sleep spindle (SS) pattern and integrate the pattern as a disturbance into the LDS. The control theory part is used to *design* a new residual which suppresses SS signals but faithfully represents other errors generated by the LDS model. Thus by selectively suppressing SS pattern, the objectives of the OSAD problem are achieved.

For example, consider Figure 1. The top frame shows a typical EEG time series with both the SS and KC highlighted. The middle frame shows a typical residual time series based on an LDS model. The bottom frame shows a new residual designed to solve the OSAD problem. Notice that the error due to the presence of SS is suppressed but the residual due to the appearance of KC remains unaffected.

The main contributions of the paper are:

- We introduce the Online Selective Anomaly Detection(OSAD) to address the requirement of selectively reporting sleep anomalies based on specifications by domain experts.
- In order to solve OSAD, we combine techniques from data mining and control theory. In particular we will use a Linear Dynamic System (LDS) to model the underlying data generating process and use control theory techniques to design an appropriate residual system.

The rest of the paper is as follows. In Section 2, we rigorously define the OSAD problem. In Section 3 we present our methodology to infer the parameters of the LDS and use control theory to design a new residual system. In Section 4, we apply our approach to real sleep data and evaluate our results. We overview related work in Section 5 and conclude in Section 6 with a summary and potential ideas for future research.

2. PROBLEM DEFINITION

In this section we present our problem statement for selective anomaly detection.

The starting point is an observed time series of N points $y = \{y_i\}_{i=1}^N$ where each $y_i \in \mathbb{R}^m$. Furthermore, we assume that the y measures the output of a system which is generated from a latent variable $x \in \mathbb{R}^n$. The relationship between x and y is governed by a standard Linear Dynamic

system (LDS) model [25] which is specified as

$$\begin{aligned} x(t+1) &= \mathbf{A}x(t) \\ y(t) &= \mathbf{C}x(t) \end{aligned}$$

Here \mathbf{A} is an $n \times n$ state matrix which governs the dynamics of the LDS while \mathbf{C} is an $m \times n$ observation matrix. The modern convention is to represent the LDS as graphical model as shown in Figure 2. The state of the system, x , evolves according to LDS beginning at time $t = 0$, with value x_0 . The standard learning problem is as follows.

PROBLEM 1 (LEARNING PROBLEM). *Given an observable time series $\{y_i\}_{i=1}^N$ and assuming that the observed y and the latent x are governed by an LDS, infer \mathbf{A} and \mathbf{C} .*

The standard LDS inference problem has been extensively studied in both the machine learning and control theory literature. Several algorithms have been proposed including those based on gradient descent, Expectation Maximization, subspace identification and spectral approaches [26, 28, 16, 3]. Several extensions of LDS to include non-linear relationships as well as to include stochastic disturbances have been proposed. However, for sleep analysis, the above LDS will suffice. For the sake of completeness, in the Appendix we will describe a simple but effective approach for inferring \mathbf{A} and \mathbf{C} based on a spectral method [3].

The standard approach to detect outliers using an LDS is to use the inferred \mathbf{A} and \mathbf{C} matrices to compute the latent and observed error variables as:

$$\begin{aligned} \epsilon(t) &:= x(t) - \hat{x}(t) \\ e(t) &:= y(t) - \hat{y}(t) \end{aligned}$$

where \hat{x} and \hat{y} are estimated using LDS. Then given a threshold parameter δ , an anomaly is reported whenever, $e(t) > \delta$. However, our objective is not to report all anomalies but suppress some known user-defined patterns or even known anomalous pattern. We now formalize the notion of pattern.

DEFINITION 1. *A pattern \mathbf{P} is a user-defined matrix which operates in the latent space.*

In our context, we will design a specific matrix \mathbf{P} for a sleep spindle. The matrix \mathbf{P} is integrated into the LDS as

$$\begin{aligned} x(t+1) &= \mathbf{A}x(t) + \mathbf{P}\zeta(t) \\ y(t) &= \mathbf{C}x(t) \end{aligned}$$

We are now ready to define the design part of the OSAD problem.

PROBLEM 2 (DESIGN PROBLEM). *Given an LDS, a pattern \mathbf{P} in the latent space, design a residual $r(t)$ such that*

$$r(t) = \begin{cases} 0 & \text{if } \epsilon(t) = \mathbf{P}\zeta(t) \\ \mathbf{S}e(t) & \text{otherwise} \end{cases}$$

Here \mathbf{S} is suitably defined linear transformation on $e(t)$. Notice that the residual $r(t)$ depends both on the latent error $\epsilon(t)$ and the observed error $e(t)$. In practice, $r(t)$ will never be exactly zero when the pattern \mathbf{P} is active but will have small absolute values.

3. THE OSAD METHOD

In this section we propose a method based on statistical inference and control theory to provide a solution of the




Figure 2: A linear dynamic system is a model which defines a linear relationship between the latent (or hidden) state of the model and observed outputs. The LDS parameters \mathbf{A} and \mathbf{C} need to be estimated from data. The LDS can also be used to model the relationship between the latent and the observed residuals (right figure).

OSAD problem. Using the LDS, we first develop a **Dynamic Residue Model (DRM)**. Then we will show how to adjust the DRM parameters in order to design a residual $r(t)$ which will satisfy the constraints of the problem, i.e. the selected anomalous pattern will be canceled (or projected out) in the generated residual space.

3.1 DRM Formulation

Assume data is generated by an LDS. Any deviation of the state from its expected value can be captured by a structured error model. Intuitively, the discrepancy between the observed error $e(t)$ and latent error $\varepsilon(t)$ is modeled by the same LDS (because of linearity):

$$\begin{aligned}\varepsilon(t+1) &= \mathbf{A}\varepsilon(t) + \mathbf{P}\xi(t) \\ e(t) &= \mathbf{C}\varepsilon(t)\end{aligned}$$

The above error model can be used to detect changes occurring in the latent space.

We design a feedback loop (as shown in Figure 3) to effect the output of the error model. In particular a function of the residual will be used to manipulate the changes in the error. The design objective will be to map the anomalies generated by the \mathbf{P} pattern into the null space of the new residual. The DRM based on this feedback design is developed as follows:

To design the feedback we define two transformation matrices \mathbf{W} and \mathbf{F} for error values to be weighted as:

$$\begin{aligned}r(t) &:= \mathbf{W}e(t) \\ u(t) &:= \mathbf{F}e(t)\end{aligned}$$

\mathbf{F} will be used as the feedback gain matrix and maps the error to the feedback vector $u(t)$, and \mathbf{W} is the residual weighting matrix that generates the new residual $r(t)$. Now feeding back $u(t)$ into the LDS (as shown in Figure 2), with $u(t)$, the residual dynamic model will be:

$$\begin{aligned}\hat{x}(t+1) &= \mathbf{A}\hat{x}(t) + u(t) \\ &= \mathbf{A}\hat{x}(t) + \mathbf{F}e(t) \\ &= \mathbf{A}\hat{x}(t) + \mathbf{F}(\mathbf{y}(t) - \hat{\mathbf{y}}(t)) \\ &= \mathbf{A}\hat{x}(t) + \mathbf{F}(\mathbf{C}x(t) - \mathbf{C}\hat{x}(t)) \\ &= \mathbf{A}\hat{x}(t) + \mathbf{F}\mathbf{C}x(t) - \mathbf{F}\mathbf{C}\hat{x}(t) \\ &= (\mathbf{A} - \mathbf{F}\mathbf{C})\hat{x}(t) + \mathbf{F}\mathbf{C}x(t) \\ &= (\mathbf{A} - \mathbf{F}\mathbf{C})\hat{x}(t) + \mathbf{F}y(t)\end{aligned}$$

Notice that since the residual is a linear transformation

of the error, its rank (suppose $r(t) \in \mathbb{R}^p$) can not be larger than the observation dimension, i.e., $p \leq m$.

We are now able to define the dynamic of the latent error as:

$$\begin{aligned}\varepsilon(t+1) &= x(t+1) - \hat{x}(t+1) \\ &= \mathbf{A}x(t) - (\mathbf{A} - \mathbf{F}\mathbf{C})\hat{x}(t) - \mathbf{F}y(t) \\ &= \mathbf{A}x(t) - \mathbf{A}\hat{x}(t) - \mathbf{F}\mathbf{C}\hat{x}(t) + \mathbf{F}\mathbf{C}x(t) \\ &= (\mathbf{A} - \mathbf{F}\mathbf{C})(x(t) - \hat{x}(t)) \\ &= (\mathbf{A} - \mathbf{F}\mathbf{C})\varepsilon(t)\end{aligned}$$

and the residue $r(t)$ is obtained as:

$$\begin{aligned}r(t) &= \mathbf{W}(y(t) - \hat{y}(t)) \\ &= \mathbf{W}(\mathbf{C}x(t) - \mathbf{C}\hat{x}(t)) \\ &= \mathbf{W}\mathbf{C}(x(t) - \hat{x}(t)) \\ &= \mathbf{W}\mathbf{C}\varepsilon(t)\end{aligned}$$

We therefore have the following dynamic model for the latent error:

$$\begin{aligned}\varepsilon(t+1) &= (\mathbf{A} - \mathbf{F}\mathbf{C})\varepsilon(t) \\ r(t) &= \mathbf{W}\mathbf{C}\varepsilon(t)\end{aligned}$$

Notice that the observed residue $r(t)$ is governed by state error $\varepsilon(t)$ through matrix $\mathbf{W}\mathbf{C}$ while it evolves in time through $\mathbf{A} - \mathbf{F}\mathbf{C}$.

To simplify the notation, denote $\mathbf{C}_f = \mathbf{W}\mathbf{C}$ and $\mathbf{A}_f = \mathbf{A} - \mathbf{F}\mathbf{C}$. The DRM is then defined as:

$$\begin{aligned}\varepsilon(t+1) &= \mathbf{A}_f\varepsilon(t) \\ r(t) &= \mathbf{C}_f\varepsilon(t)\end{aligned}$$

The graphical diagram for this error model is shown in Figure 3.

3.2 OSAD Parameter Design

In this section we address the problem of designing the \mathbf{F} and \mathbf{W} matrix with objective of making the DRM insensitive to anomalies generated by \mathbf{P} . The overarching design is shown in Figure 4 and is related to the use of control theory for fault diagnosis [22, 23, 5]. A typical LDS model will output the observed error $e(t)$. However, the OSAD model has a feedback loop which takes \mathbf{W} and \mathbf{F} matrices as input and return a variable $u(t)$ which is fed back into the model. The observed error is also transformed by a \mathbf{W} matrix. The \mathbf{F} and the \mathbf{W} matrices satisfy the constraints which involve the \mathbf{A} , \mathbf{C} and the \mathbf{P} matrices.

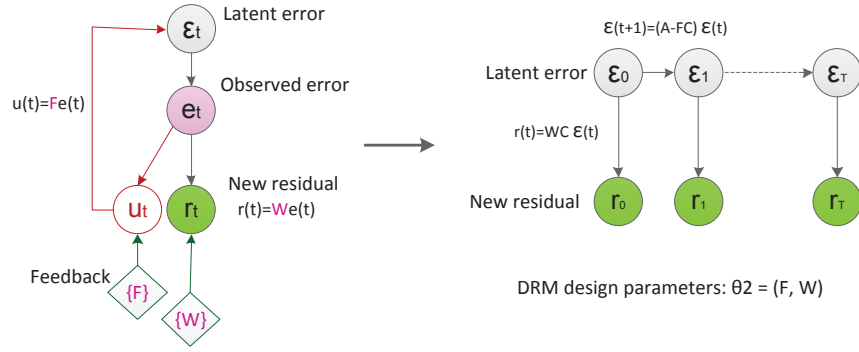


Figure 3: Using parameter \mathbf{F} a virtual input $u(t)$ is generated to feed the error back to the latent space. The error $e(t)$ is then calibrated by \mathbf{W} to generate a new residual space $r(t)$.

Since the model is time-dependent, we follow a standard approach and map the model into the frequency domain using a \mathcal{Z} -transform to design the \mathbf{W} and \mathbf{F} matrices. In the frequency domain, it will be easier to design matrices \mathbf{W} and \mathbf{F} such that $\mathbf{WC}(\mathbf{A} - \mathbf{FC}) = 0$ and $\mathbf{WCP} = 0$.

DEFINITION 2. The Z-transform of a discrete-time sequence $x(k)$ is the series $X(z)$ defined as

$$X(z) = \mathcal{Z}\{x(k)\} = \sum_{k=0}^{\infty} x(k)z^{-k}.$$

OBSERVATION 1. Two important (and well known) properties of the Z-transform are linearity and time shifting:

$$ax(k) + by(t) \xrightarrow{\mathcal{Z}} aX(z) + bY(z)$$

$$x(k+b) \xrightarrow{\mathcal{Z}} z^b X(z)$$

Applying Z-transform $\mathcal{Z}()$ to the DRM yields:

$$\begin{aligned} z\mathcal{E}(z) &= \mathbf{A}_f\mathcal{E}(z) + \mathbf{P}\xi(z) \\ \mathcal{E}(z) &= (z\mathbf{I} - \mathbf{A}_f)^{-1}\mathbf{P}\xi(z) \end{aligned}$$

and:

$$\begin{aligned} R(z) &= \mathbf{C}_f\mathcal{E}(z) \\ &= [\mathbf{C}_f(z\mathbf{I} - \mathbf{A}_f)^{-1}\mathbf{P}]\xi(z) \end{aligned}$$

in which $\xi(z) = \mathcal{Z}(\xi(t))$, $\vartheta = \mathcal{Z}(\vartheta(t))$, $R(z) = \mathcal{Z}(r(t))$. The transfer gain between ξ and R :

$$\mathbf{G}_\xi(z) := \mathbf{C}_f(z\mathbf{I} - \mathbf{A}_f)^{-1}\mathbf{P}$$

Thus if \mathbf{G}_ξ would be zero, the residual $R(z)$ is independent of the $\xi(z)$. In the other word, to make $R(z)$ independent of $\xi(z)$, one must null the space of $\mathbf{G}_\xi(z)$. Then whenever \mathcal{P} occurs it is transferred by a zero gain to the residual space. To find the null space $\mathbf{G}_\xi(z) = 0$, we expand it as:

$$\begin{aligned} \mathbf{G}_\xi(z) &= z^{-1}\mathbf{C}_f(\mathbf{I} + \mathbf{A}_f z^{-1} + \mathbf{A}_f^2 z^{-2} + \dots)\mathbf{P} \\ &= 0 \end{aligned}$$

The sufficient conditions for $\mathbf{G}_\xi(z)$ to be nulled are $\mathbf{C}_f\mathbf{P} = 0$ and either $\mathbf{C}_f\mathbf{A}_f = 0$ or $\mathbf{A}_f\mathbf{P} = 0$. Thus we have the following result.

THEOREM 1. For a DRM, a sufficient condition for $\mathbf{G}_\xi(z) = 0$ is

$$\mathbf{C}_f\mathbf{P} = 0 \text{ and } \{\mathbf{C}_f\mathbf{A}_f = 0 \text{ or } \mathbf{A}_f\mathbf{P} = 0\}$$

Now as $\mathbf{C}_f = \mathbf{WC}$, for $\mathbf{C}_f\mathbf{P} = 0$ it is sufficient that \mathbf{WC} be orthogonal to \mathbf{P} . Furthermore for $\mathbf{C}_f\mathbf{A}_f = 0$, it is sufficient to design a matrix \mathbf{A}_f such that its left eigenvectors corresponding to the zero eigenvalue are orthogonal to \mathbf{P} . Similarly, for $\mathbf{A}_f\mathbf{P} = 0$, it is sufficient to design a matrix \mathbf{A}_f , such that the right eigenvectors corresponding to the zero eigenvalues are orthogonal to \mathbf{P} . See Appendix A.

Now, it design a system which operates in an online fashion we proceed as follows. From the definition of residue:

$$r(t) = \mathbf{W}[y(t) - \hat{y}(t)]$$

Using the Z-transform, the computational form of the residual will be:

$$R(z) = [\mathbf{W} - \mathbf{C}_f(z\mathbf{I} - \mathbf{A}_f)^{-1}\mathbf{F}]Y(z)$$

Since $\mathbf{C}_f\mathbf{A}_f = 0$:

$$\mathbf{C}_f(z\mathbf{I} - \mathbf{A}_f)^{-1}\mathbf{F} = z^{-1}\mathbf{C}_f\mathbf{F}$$

Replacing this result to the above $R(z)$ equation:

$$R(z) = (\mathbf{W} - z^{-1}\mathbf{C}_f\mathbf{F})Y(z)$$

Applying the inverse Z-transform, the equation will be:

$$r(t) = [\mathbf{W} - \mathbf{C}_f\mathbf{F}] \begin{bmatrix} y(t) \\ y(t-1) \end{bmatrix}$$

This clearly says that the residual can be represented directly in terms of the observations. This property is crucial to make the anomaly detection system operate in near real-time.

3.3 Eigenpair Assignment and the \mathbf{F} Matrix

In this section we explain the eigenpair assignment problem and its solution which is used for designing the matrix \mathbf{F} . Recall from Theorem 1, that we require either $\mathbf{C}_f\mathbf{A}_f = 0$ or $\mathbf{A}_f\mathbf{P} = 0$.

PROBLEM 3. Given a set of scalars $\{\lambda_i\}$ and a set of n -vectors $\{v_i\}$ (for $i = 1, 2, \dots, n$), find a real matrix \mathbf{A}_o ($m \times n$) such that the eigenvalues of \mathbf{A}_o are precisely those of the set of scalars $\{\lambda_i\}$ with corresponding eigenvectors the set $\{v_i\}$.

Given the residue model transition matrix $\mathbf{A}_f = \mathbf{A} - \mathbf{FC}$, the problem is to find a matrix \mathbf{F} such that this matrix has the eigenvalues $\{\lambda_i\}$ corresponding to eigenvectors $\{v_i\}$, i.e.,:

$$(\mathbf{A} - \mathbf{FC})v_i = \lambda_i v_i$$

We will now give a concrete example. Assume we have an LDS system given as

$$\begin{aligned}\varepsilon(t+1) &= \mathbf{A}\varepsilon(t) + \mathbf{P}\xi(t) \\ e(t) &= \mathbf{C}\varepsilon(t)\end{aligned}$$

Assume we have identified the \mathbf{A} and \mathbf{C} matrices as

$$\mathbf{A} = \begin{pmatrix} 0.5 & 0.3 \\ 0.3 & 0.2 \end{pmatrix} \text{ and } \mathbf{C} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \text{ and } \mathbf{P} = \begin{pmatrix} 1 & 1 \\ 2 & 2 \end{pmatrix}$$

Now, to form the OSAD model, we have to identify \mathbf{W} and \mathbf{F} such that:

1. \mathbf{W} is in the null space of $\mathbf{C}\mathbf{P}$ and
2. $\mathbf{A} - \mathbf{F}\mathbf{C}$ has its left eigenvectors (corresponding to the 0 eigenvalue), the rows of $\mathbf{W}\mathbf{C}$.

Since \mathbf{C} is the identity matrix, an example of \mathbf{W} is

$$\mathbf{W} = \begin{pmatrix} 2 & -1 \\ 2 & -1 \end{pmatrix}$$

Similarly, an example of \mathbf{F} matrix is

$$\mathbf{F} = \begin{pmatrix} 0.0 & 0.2 \\ -0.7 & 0 \end{pmatrix}$$

As mentioned, the residual matrix is given by

$$r(t) = \begin{pmatrix} 1.3 & -1.4 \\ 1.3 & -1.4 \end{pmatrix} \begin{pmatrix} y(t) \\ y(t-1) \end{pmatrix}$$

4. EXPERIMENTAL RESULT

We now report on the experiments that have been carried out to test the effective of the proposed OSAD solution on sleep data. Our particular focus will be determining if OSAD can recognize sleep spindle and K-Complex anomalies and selectively raise an alert for non-Spindle anomalies.

4.1 Sleep Data Set

Our data set consists of EEG time series from four health controls (age 25-36) as described in [10]. Recordings were made with an Alice-4 system (Respironics, Murraysville PA, USA) at the Woolcock Institute of Medical Research, at Sydney University, using 6 EEG channels with a sampling rate of 200 Hz, and electrodes positioned according to the International 10-20 system [19, 4], see Figure 5. In this study we only examine the Cz electrode. A notch filter at 50 Hz (as provided by the Alice-4 system) was used to remove mains voltage interference. No other hardware filters were used. Spindles and K-Complexes were labeled using another automation program and then manually evaluated. As previously noted, while data from only four subjects were used, a typical EEG session generates a large amount of personal data.

4.2 Inference of \mathbf{A} and \mathbf{C} Matrices

Our first task is to learn the \mathbf{A} and \mathbf{C} matrices from the LDS for each subject. Others have reported, and our experiments confirm, that EEG of each subject tends to different and separate models need to learnt per subject. For each subject we took a sample of size 2000 (10 seconds) of EEG time series which did not contain either sleep spindle or K-Complex. We then formed a 2000×6 data matrix, \mathbf{O} . The columns of the \mathbf{O} matrix are time series associated with the




Figure 5: The position of scalp electrodes for EEG experiment follows the International 10-20 system [19, 4].

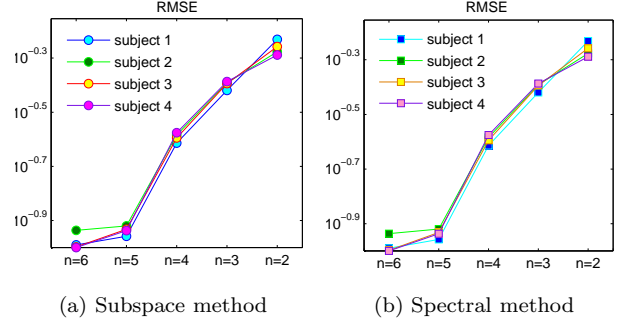


Figure 6: The RMSE error obtained from both methods are comparable. Notice the RMSE increases as the rank of LDS is reduced.

six channels of EEG. We used both subspace and spectral methods to infer the matrices \mathbf{A} and \mathbf{C} . Both these methods are based on SVD decomposition of the \mathbf{O} matrix and require as input the rank required of the inferred matrices. We evaluated the inferred matrices using RMSE and the results are shown in Figure 6a and Figure 6b. Both the subspace and spectral methods have similar performance and RMSE goes up significantly when the rank falls below five. We selected a rank six matrix (maximum possible rank) for both \mathbf{A} and \mathbf{C} . In terms of running time, the two methods are comparable as we have to carry out an SVD of a relatively small 6×6 matrix.

4.3 Detection of SS and K-Complex

For each of the four subjects, statistics of the labeled sleep spindles and K-Complexes and those detected by the LDS are shown in Table 2. For LDS detection, we used a threshold derived from CUSUM which automatically adjusts for mean and standard deviation of the observed residual time series $e(t)$. To specify a CUSUM threshold we applied the alpha and beta approach in [17] and we set the probabilities of a false positive and a false negative to 10^{-4} and the change detection parameter to 1 sigma, in all subjects.

In all four subjects, the LDS residual slightly under predicts the number of spindles and K-Complexes. Since each labeled and predicted SS and K-Complex spans a time-interval, we have modified the definitions of **precision** and **recall** to

Table 2: Summary statistics of results. LDS is quite accurate but tends to over-predict the number of anomalies.

	No. of Labeled Anomalies		No. of Detected Anomalies	
	Spindle	K-Complex	Spindle	K-Complex
subject 1	170	277	164	251
subject 2	6	13	6	11
subject 3	23	38	21	37
subject 4	141	205	132	186

Table 3: Summary statistics for spindles. LDS has higher precision than recall and total length of predicted interval is higher than the length of labeled intervals.

	Total time of spindles		Performance	
	Labeled in min	Detected in min	Recall	Precision
subject 1	129.8	168.74	71.24%	95.53%
subject 2	3.45	3.55	97.18%	97.38%
subject 3	15.15	16.23	83.88%	95.66%
subject 4	93.5	103.2	79.15%	95.42%

Table 4: Summary statistics for K-Complex. Both precision and recall are high. Total length of predicted interval is higher than labeled intervals.

	Total time of K-Complex		Performance	
	Labeled in min	Detected in min	Recall	Precision
subject 1	198.23	216.35	90.45%	93.43%
subject 2	11.48	11.25	92.76%	94.12%
subject 3	21.39	24.56	91.01%	92.06%
subject 4	147.68	160.49	91.28%	93.73%

account for the intervals. For a given subject, let $\{[a_i, b_i]\}_{i=1}^n$ be the intervals of the labeled anomalies (spindles or K-Complex). Let $\{[a'_j, b'_j]\}_{j=1}^m$ be the predicted spindles. Then

$$\text{precision} = \frac{\sum_{i=1}^n \sum_{j=1}^m |[a_i, b_i] \cap [a'_j, b'_j]|}{\sum_{j=1}^m |[a'_j, b'_j]|}$$

and

$$\text{recall} = \frac{\sum_{i=1}^n \sum_{j=1}^m |[a_i, b_i] \cap [a'_j, b'_j]|}{\sum_{j=1}^m |[a'_j, b'_j]|}$$

Here, $|[a_i, b_i]|$, is the number of points in the time interval $[a_i, b_i]$. With these definitions in place, Table 3 and Table 4 show the precision and recall SS and K-Complex across all the subjects. In general both precision and recall are high across subjects, but precision is significantly more higher than recall. For SS, the recall varies more than precision ranging for 71.24% to 97.18%. Also notice that the length of detection of both SS and K-Complex is higher compared to their labeled lengths.

4.4 Evaluation across Subjects

We now investigate the transfer properties of the inferred LDS across subjects. That is, we learn the **A** and **C** matrices on one subject and evaluate it against another. We just

Table 5: Recall across subjects. A substantial reduction in accuracy when model of one subject is evaluated against the EEG of another.

	A_1, C_1, W_1, F_1	A_2, C_2, W_2, F_2	A_3, C_3, W_3, F_3	A_4, C_4, W_4, F_4
subject 1	71.24%	38.13%	44.13%	41.29%
subject 2	41.32%	97.18%	35.26%	37.85%
subject 3	48.74%	43.21%	83.88%	44.43%
subject 4	51.26%	43.81%	35.36%	79.15%

Table 6: Precision across the subjects. Again, a substantial reduction in accuracy when model of one subject is evaluated against another.

	A_1, C_1, W_1, F_1	A_2, C_2, W_2, F_2	A_3, C_3, W_3, F_3	A_4, C_4, W_4, F_4
subject 1	95.53%	41.11%	47.19%	43.67%
subject 2	39.54%	97.38%	37.82%	39.21%
subject 3	48.21%	41.29%	95.77%	41.83%
subject 4	51.77%	53.34%	33.49%	95.42%

Table 7: Recall and Precision on each subject evaluated against an averaged model. Again, a substantial reduction in accuracy compared to individual models.

	<i>Recall</i>	<i>Precision</i>
subject 1	69.35%	51.39%
subject 2	65.43%	57.22%
subject 3	61.77%	61.47%
subject 4	68.12%	53.92%

focus on the anomaly. The **recall** and **precision** results are shown in Table 5 and Table 6 respectively. The diagonal of the table corresponds to the results in Table 3 and Table 4. It is clear that there is a substantial reduction in accuracy and that indeed the EEG of subjects varies substantially. We have also computed the "average" **A** and **C** matrix and evaluated against all the four subjects. The results are shown in Table 7. While there is an improvement compared to results in Table 5 and Table 6, the absolute performance is still quite low compared to the situation where the learning was customized per individual subject.

4.5 Performance of Designed Residual

In this section we evaluate whether the new residual $r(t)$ satisfies the design criterion. Recall, $r(t)$ was designed to suppress the signal whenever a sleep spindle (SS) appears and behave like the observed error $e(t)$ in otherwise. Figure 7 shows the distribution for $|r(t) - e(t)|_2$ for values of t when t is in (and not in) the predicted SS interval $[a'_j, b'_j]$ for some j . It is clear that the distribution when t is in a predicted SS interval is towards the right compared to when it is not in the interval. This is because in an SS interval, $r(t)$ will have a small absolute value (by design). In a non-SS interval, $r(t)$ will be a linear function of $e(t)$, as $r(t) = \mathbf{W}e(t)$. This behavior is observed across subjects suggesting that in all cases that $r(t)$ is behaving as designed. Furthermore in Figure 8, we plot the $|r(t)|$ against $|e(t)|$ when t is not in a




Figure 7: Comparison of the distribution of the norm of $r(t) - e(t)$ for SS and non-SS intervals. In all four subjects the designed residual suppresses spindles as designed as the norm is higher for SS intervals.

spindle interval. Again we observe a straight line behavior, providing further confirmation that $r(t)$ is behaving according to specifications.




Figure 8: the $|r(t)|$ against $|e(t)|$ when t is not in a spindle interval as $r(t) = \mathbf{W}e(t)$

4.6 Delay in Detection of Anomalies

OSAD detects anomalies in near real time. We now discuss the lag between the appearance of a SS and before it is reported by the LDS. Figure 9 presents the delay distributions for subject 1 and subject 4 who experienced 164 and 132 labeled sleep spindles, respectively. In general, the predicted SS interval are longer and contain the actual intervals. This is confirmed in Figure 10 which shows one specific example of the location of the labeled sleep spindle and the predicted interval. In this case (which is typical), the prediction of SS begins before and ends later than the labeled spindle. Table 8 shows the results of the mean delay between matched intervals. Thus a mean of (a_i, a'_i) equal to -0.0678 implies that on average, there was a delay of $1/200$

Table 8: Delay statistics. The lag between appearance and prediction of SS is, on average, a fraction of a second.

	$\text{Mean}(a_i - a'_i)$	$\text{Mean}(b_i - b'_i)$	$\text{Std}(a_i - a'_i)$	$\text{Std}(b_i - b'_i)$
subject 1	- 0.0678	- 0.0961	0.0589	0.0995
subject 2	0.0041	0.0016	0.0113	0.0089
subject 3	-0.0426	0.0663	0.0550	0.0478
subject 4	-0.0340	-0.0480	0.0474	0.0407

second before LDS reported an anomaly. On the other hand for subject 2 there the SS was, on average, reported before it showed up in the labeled sequence. As noted in [10], this is consistent with the observation (and confirmed by double-blind scoring) that the labeling of SS is more conservative i.e., SS are labeled for a shorter duration than what they should be.

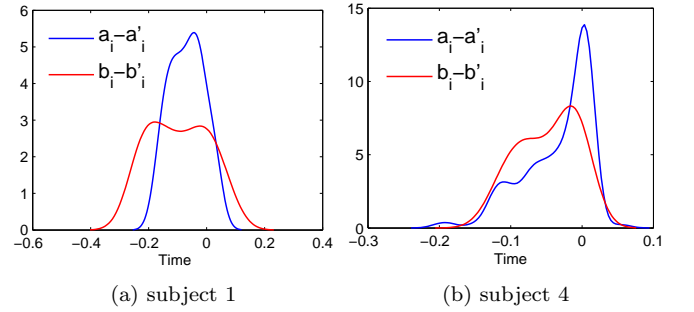


Figure 9: OSAD provides near real time detection. The delay between the actual appearance of a spindle and the predicted appearance is a fraction of a second. Similarly the lag between when the actual spindle disappears and it is reported to disappear is very small too. The x-axis is in seconds.

5. RELATED WORK

Automatic detection of sleep spindles is now an important topic in biomedical research. Different techniques including FFTs, wavelet analysis and autoregressive time series modeling have been applied for sleep spindle detection [24, 11, 13]. Attempts to integrate SVM to detect sleep spindles have also been explored [1]. There seems to be a large variability between sleep EEG across subjects. In our experiments we have also observed this phenomenon. This combined with the large amount of EEG noise has resulted in low level of agreement on the exact profile of sleep spindle [20].

The use of Linear Dynamical Systems (LDS) to model time series is ubiquitous both in computer science [25] and control theory [28, 16, 18, 7, 15]. Expressing LDS in the language of graphical models and connections with HMM have been extensively examined in machine learning. The use of LDS for anomaly detection has also been investigated in network anomaly detection, among other areas [27]. The use of subspace identification methods for inferring the parameters of LDS have been discussed by Overschee [28]. Subspace methods estimate LDS parameters through a spectral decomposition of a matrix of observations to yield an estimate of the underlying state space. Subspace methods have low computational cost, are robust to perturbations and are relatively




Figure 10: Top: Cz data and a typical sleep spindle labeled. Bottom: Residual and detected sleep spindle. In general the predicted spindle interval is longer than the labeled interval. The predicted interval tends to include the labeled interval, i.e., it begins earlier and finishes later. The EEG shows that the labeled intervals are actually quite conservative.

easy to implement. The recently introduced spectral learning methods are variations of the subspace method [3, 12]

The use of eigenstructure assignment to alter the residual of an LDS has been investigated in the control theory literature especially in the context of fault diagnosis [2]. Our approach closely follows the work Patton et. al. [21] who have used eigenstructure assignment for altering the LDS model using feedback. Other variations of LDS and fault diagnosis are discussed in [6, 22, 5, 23].

6. CONCLUSION

In this paper we have introduced a new problem, the On-line Selective Anomaly Detection (OSAD) to capture a specific scenario in sleep research. Scientists working on sleep EEG data required an alert system, which trigger alerts on selected anomalies. For example, sleep stage two is characterized by two known anomalies: sleep spindle and K-complex. The requirement was to design a system which detected both anomalies but only generated an alert when a non sleep spindle anomaly appeared. We combined methods from data mining, machine learning and control theory to design such a system. Experiments on real data set demonstrate that our approach is accurate and produces the required results and is potentially applicable to many other situations. We also note that data from sleep EEG provides a fertile ground to apply existing data mining methodologies and potentially design new computational problems and algorithms.

7. ACKNOWLEDGMENTS

This work is partially supported by NICTA¹. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

¹<http://nicta.com.au/>

8. REFERENCES

- [1] N. Acir and C. Güzelış. Automatic recognition of sleep spindles in eeg via radial basis support vector machine based on a modified feature selection algorithm. *Neural Computing & Applications*, 14(1):56–65, 2005.
- [2] A. Andry, E. Shapiro, and J. C. Chung. Eigenstructure assignment for linear systems. *Aerospace and Electronic Systems, IEEE Transactions on*, AES-19(5):711–729, 1983.
- [3] B. Boots, S. Siddiqi, and G. Gordon. Closing the learning-planning loop with predictive state representations. In *Proceedings of Robotics: Science and Systems*, Zaragoza, Spain, June 2010.
- [4] G. Buzsaki. *Rhythms of the Brain*. Oxford University Press, 2006.
- [5] C. Chen and R. Patton. *Robust Model-Based Fault Diagnosis For Dynamic Systems*. Kluwer International Series on Asian Studies in Computer and Information Science, 3. Kluwer, 1999.
- [6] J. Chen and R. Patton. Optimal filtering and robust fault diagnosis of stochastic systems with unknown disturbances. *Control Theory and Applications, IEE Proceedings -*, 143(1):31–36, 1996.
- [7] J. H. Cochrane. *Asset Pricing*. Princeton University Press, 2001.
- [8] T. T. Dang-Vu, S. M. McKinney, O. M. Buxton, J. M. Solet, and J. M. Ellenbogen. Spontaneous brain rhythms predict sleep stability in the face of noise. *Current Biology*, 20(15):R626–R627, 2010.
- [9] S. Diekmann and J. Born. The memory function of sleep. *Nature Reviews Neuroscience*, 11(2):114–126, 2010.
- [10] A. L. DâĂŹRozario, J. W. Kim, K. K. Wong, D. J. Bartlett, N. S. Marshall, D.-J. Dijk, P. A. Robinson, and R. R. Grunstein. A new {EEG} biomarker of neurobehavioural impairment and sleepiness in sleep apnea patients and controls during extended wakefulness. *Clinical Neurophysiology*, 124(8):1605 – 1614, 2013.
- [11] F. Duman, A. Erdamar, O. Eroglu, Z. Telatar, and S. Yetkin. Efficient sleep spindle detection algorithm with decision tree. *Expert Systems with Applications*, 36(6):9980–9985, 2009.
- [12] D. Hsu, S. M. Kakade, and T. Zhang. A spectral algorithm for learning hidden markov models. *CoRR*, abs/0811.4413, 2008.
- [13] E. Huupponen, G. Gómez-Herrero, A. Saastamoinen, A. Värri, J. Hasan, and S.-L. Himanen. Development and comparison of four sleep spindle detection methods. *Artificial intelligence in medicine*, 40(3):157–170, 2007.
- [14] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME-Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- [15] G. Kitagawa. Non-gaussian state space modeling of nonstationary time series. *Journal of the American statistical association*, 82(400):1032–1041, 1987.
- [16] L. Ljung. *System Identification*. John Wiley & Sons, Inc., 2001.
- [17] D. Montgomery. *Introduction to Statistical Quality Control*. Wiley, 2004.

- [18] O. Nelles. *Nonlinear System Identification: From Classical Approaches to Neural Networks and Fuzzy Models*. Engineering Online Library. Springer, 2001.
- [19] E. Niedermeyer and F. L. da Silva. *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins, 2005.
- [20] A. Nonclercq, C. Urbain, D. Verheulpen, C. Decaestecker, P. Van Bogaert, and P. Peigneux. Sleep spindle detection through amplitude–frequency normal modelling. *Journal of neuroscience methods*, 214(2):192–203, 2013.
- [21] R. Patton and J. Chen. Robust fault detection using eigenstructure assignment: a tutorial consideration and some new results. In *Decision and Control, 1991., Proceedings of the 30th IEEE Conference on*, pages 2242–2247 vol.3, 1991.
- [22] R. Patton and J. Chen. Observer-based fault detection and isolation: Robustness and applications. *Control Engineering Practice*, 5(5):671 – 682, 1997.
- [23] R. J. Patton and J. Chen. On eigenstructure assignment for robust fault diagnosis. *International Journal of Robust and Nonlinear Control*, 10(14):1193–1208, 2000.
- [24] L. B. Ray, S. M. Fogel, C. T. Smith, and K. R. Peters. Validating an automated sleep spindle detection algorithm using an individualized approach. *Journal of sleep research*, 19(2):374–378, 2010.
- [25] S. Roweis and Z. Ghahramani. A unifying review of linear gaussian models. *Neural Comput.*, 11(2):305–345, 1999.
- [26] R. H. Shumway and D. S. Stoffer. An approach to time series smoothing and forecasting using the em algorithm. *Journal of Time Series Analysis*, 3(4):253–264, 1982.
- [27] A. Soule, K. Salamatian, and N. Taft. Combining filtering and statistical methods for anomaly detection. In *IMC '05 Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, IMC '05, pages 31–31, 2005.
- [28] P. van Overschee and L. de Moor. *Subspace identification for linear systems: theory, implementation, applications*. Number v. 1. Kluwer Academic Publishers, 1996.

APPENDIX

A. PROOF OF THEOREM 1

Theorem 1. For a DRM, a sufficient condition for $\mathbf{G}_\xi(z)$ to be zero, is

$$\mathbf{C}_f \mathbf{P} = 0 \text{ and } \{\mathbf{C}_f \mathbf{A}_f = 0 \text{ or } \mathbf{A}_f \mathbf{P} = 0\}$$

Proof: Let the set $\{\lambda_i = 0, v_i\}$, for $i = 1 : n$, be the left eigenvectors and corresponding eigenvalues of \mathbf{A}_f , i.e.

$$\begin{aligned} v_i \mathbf{A}_f &= \lambda_i v_i \\ &= 0 \end{aligned}$$

If one chooses v_1 as the rows of matrix $[\mathbf{WC}]$, then:

$$[v_1 \quad \dots \quad v_n]' \mathbf{A}_f = 0 \quad \Rightarrow \quad \mathbf{WCA}_f = 0$$

The matrix $\mathbf{A}_f = \mathbf{A} - \mathbf{FC}$, so it is sufficient to chose \mathbf{F} so that the set $\{\lambda_i = 0, v_i = [\mathbf{CW}]'\}$ to be assigned as left eigenpairs of $(\mathbf{A} - \mathbf{FC})$.

In the other side, suppose If the columns of \mathbf{P} are the right eigenvectors of \mathbf{A}_f corresponding to zero-values eigenvectors, then

$$\mathbf{A}_f v_i = 0 \quad \Rightarrow \quad \mathbf{A}_f \mathbf{P} = 0$$

So it is sufficient to chose \mathbf{F} so that the set $\{\lambda_i = 0, v_i = \mathbf{P}\}$ to be assigned as right eigenpairs of $(\mathbf{A} - \mathbf{FC})$.